



# The commissioning of CMS computing centres in the WLCG Grid

Andrea Sciabà

S. Belforte, A. Fanfani, I. Fisk, J. Flix, J. Hernández, J. Klem,  
J. Letts, N. Magini, V. Miccio, S. Padhi, P. Saiz, F. Würthwein

ACAT 2008  
3-7 November 2008  
Erice, Italy

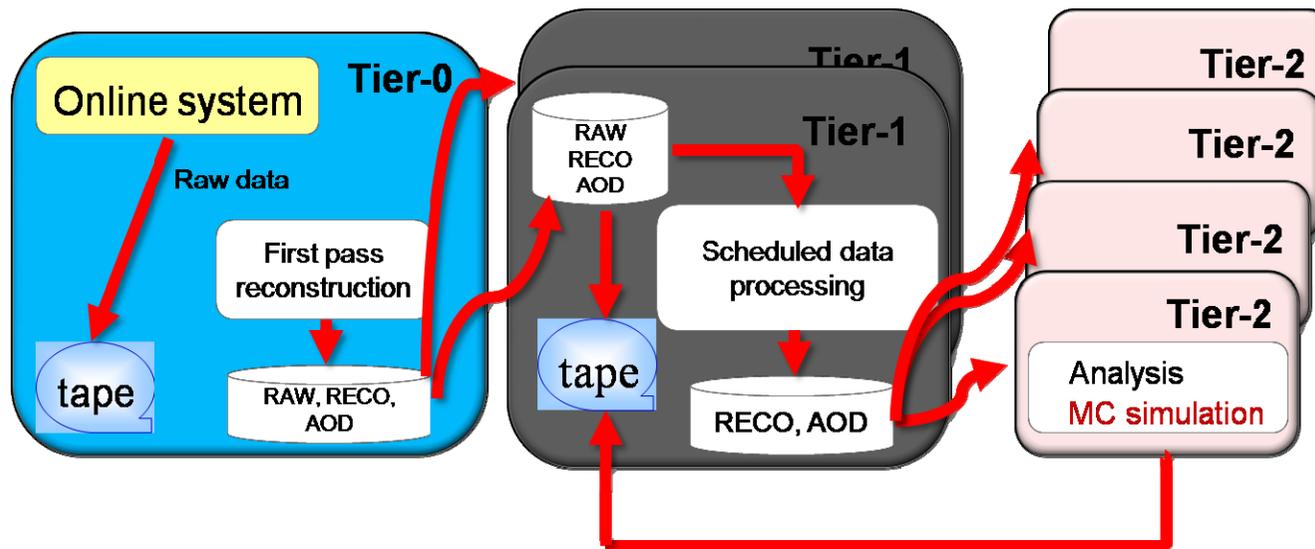


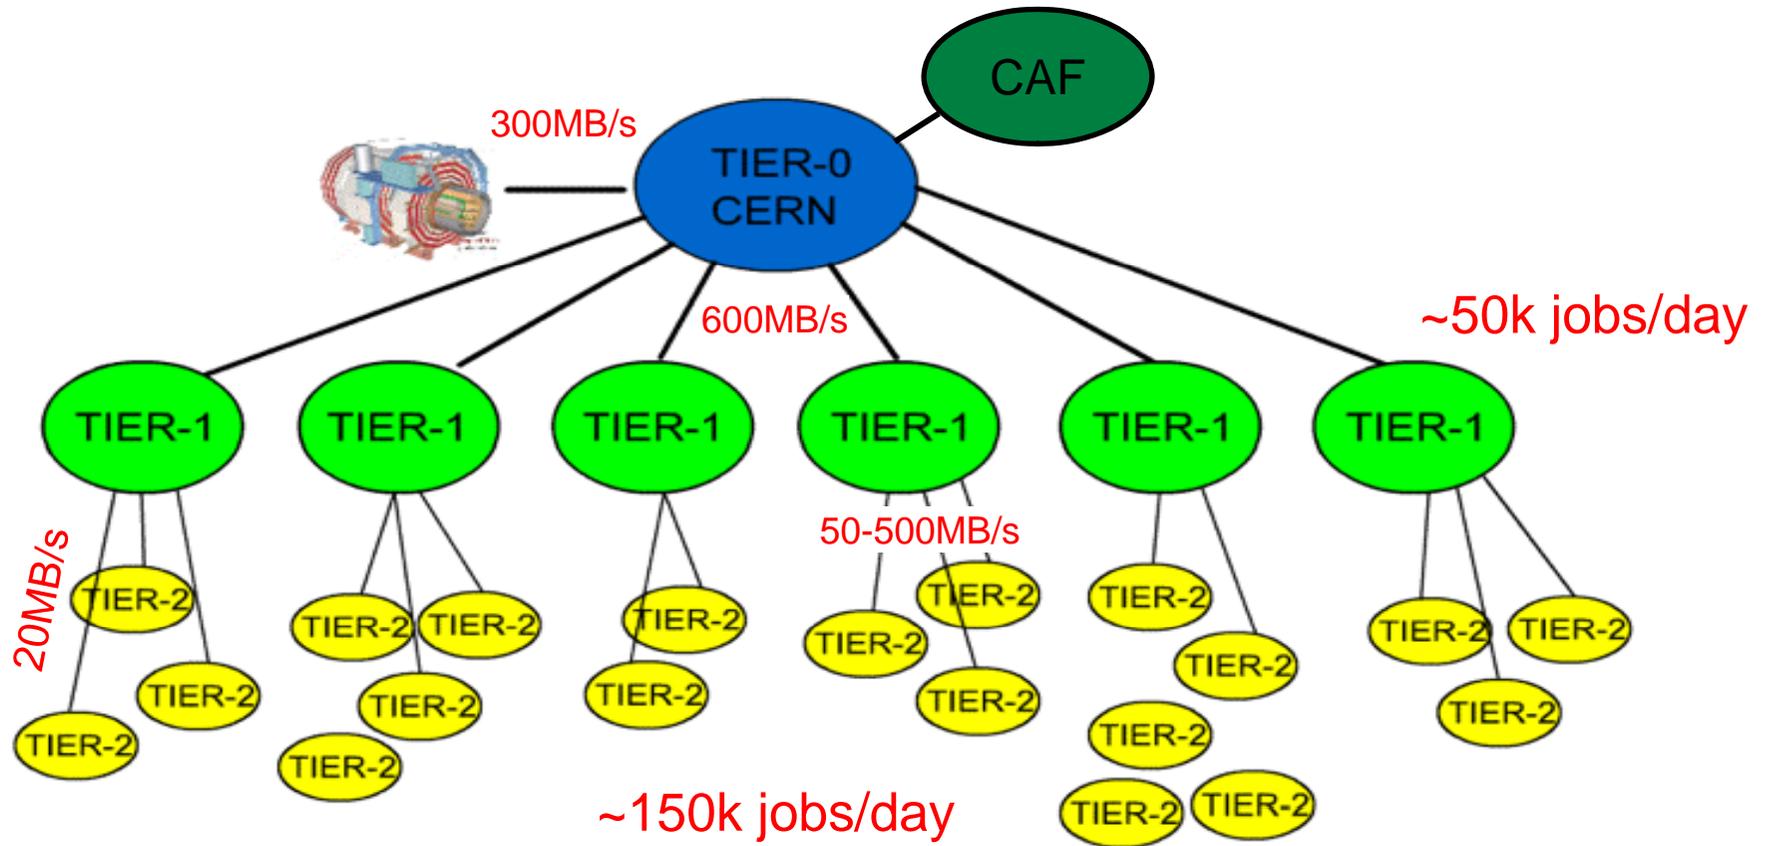
# Outline



- CMS and its computing model
- The Worldwide LHC Computing Grid
- Analysis activities during CCRC'08
- Objectives of the CMS site commissioning
- Monitoring the CMS sites
  - Site availability
  - Job Robot
  - Link commissioning
- Putting all together
  - Commissioning status
  - Site Status Board
- Conclusions

- Tier-0 (CERN)
  - Acceptance and storage of raw data
  - First pass reconstruction
  - CMS CERN Analysis Facility (latency critical data processing, high priority analysis)
- 7 Tier-1 sites
  - Archiving of raw data (shared), RECO and AOD
  - Second pass reconstruction
  - Skimming (data samples reduction)
- ~50 Tier-2 sites
  - Monte Carlo event generation
  - Data processing for physics analyses and detector studies





- Grid infrastructure to store and analyse all the LHC data (~PBs/year)
- Formed by hundreds of computing centres in > 20 countries
- Operated by three Grid projects: EGEE, OSG and NorduGrid





# Testing the Grid infrastructure



- Development phase

- Data and computing challenges have been the main instrument to evaluate reliability, scalability, functionality of
  - ☞ WLCG infrastructure: computing farms, storage systems, network links
  - ☞ Grid middleware
  - ☞ CMS software and services

- Operations phase

- The system is up and running doing real data acquisition
- The Grid infrastructure testing must be a continuous activity integrated with the monitoring system and looked at by people doing shifts



# Analysis tests during CCRC'08 (I)



- Various analysis exercises performed to measure performance and readiness of Tier-2 sites
- Site performance measurement
  - Long running CPU intensive jobs, moderate I/O
  - Long running I/O intensive jobs
  - Short running jobs with local stage out of ~10 MB file
  - Involved ~40 sites across EGEE, OSG, NorduGrid
  - Results
    - ☞ ~100'000 jobs submitted
    - ☞ Moderate I/O jobs: success rate 99.3%
    - ☞ Intensive I/O jobs: success rate 92.3%

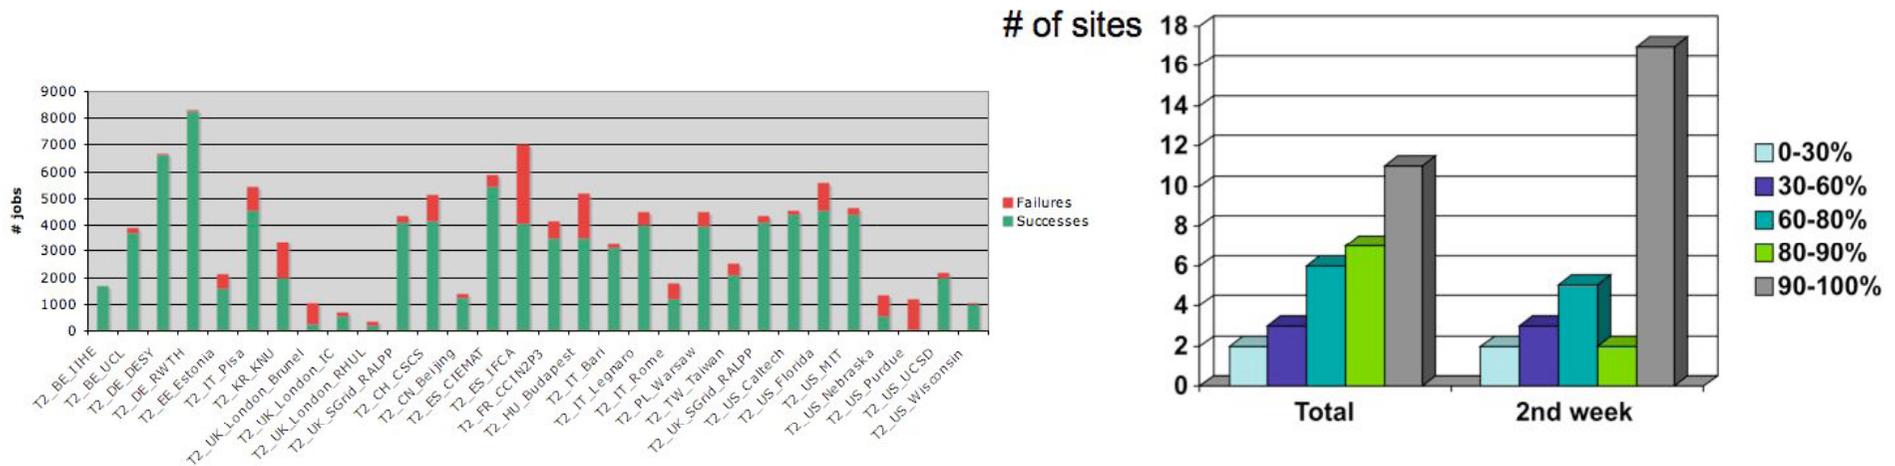
## ● Physics analysis groups

- Fake analysis groups were assigned Tier-2 sites
- Analysis-like jobs reading RECO data running for ~4 h and remote stageout of ~20 MB output

## ➤ Results

☞ 105'000 jobs submitted

☞ Initial problems with data stage out promptly fixed

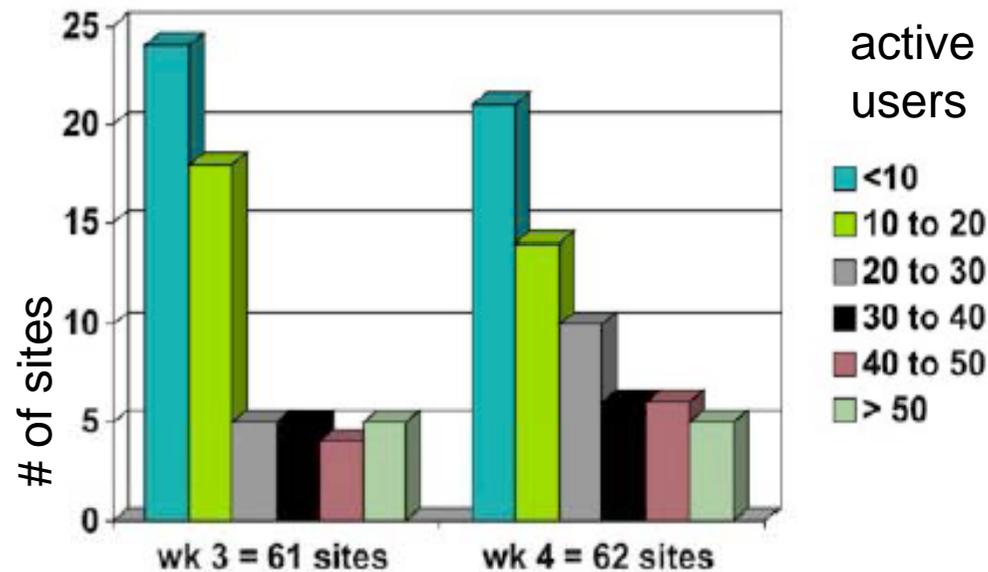


- “Chaotic” submission

- Real users were encouraged to submit analysis jobs

- Results

- ☞ High participation of users and sites



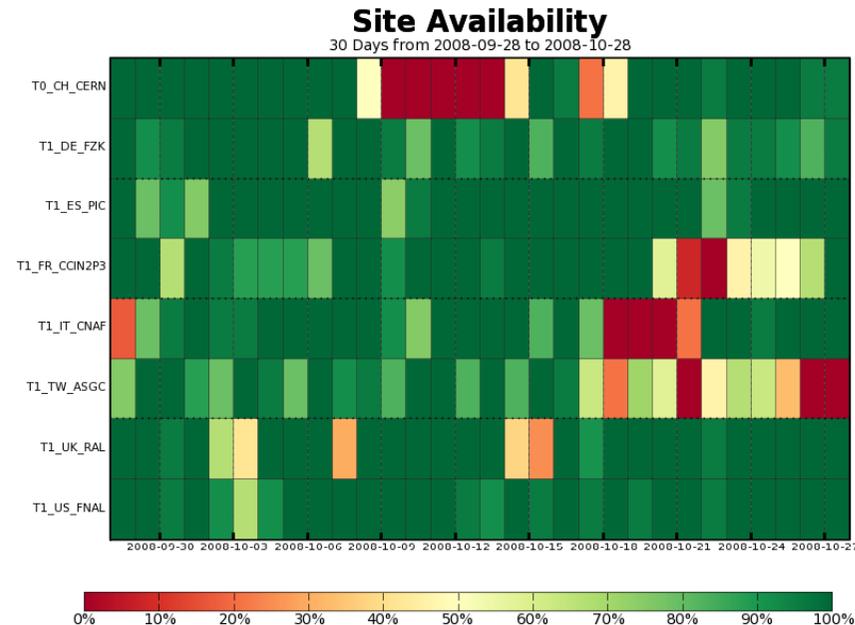


# Site Commissioning

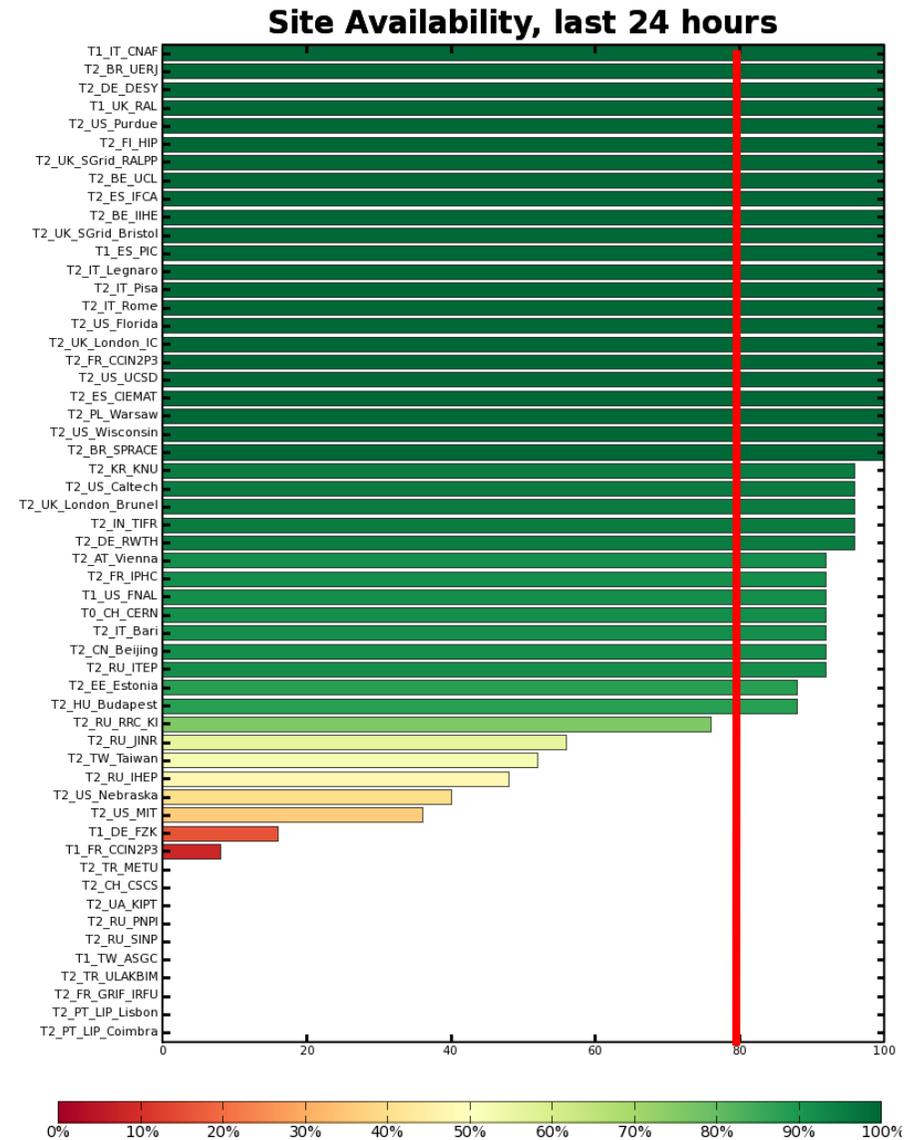


- Activity to define criteria and develop tools to
  - Measure the reliability of Grid sites when running CMS workflows
  - Verify that CMS can actually access the resources pledged
  - Provide sites with the information they need to solve eventual problems
- Intended to be run as part of the CMS computing operations
- Takes into account several sources of information
  - Site availability
  - Job Robot success rate
  - Number of commissioned links to/from other sites

- Using the Site Availability Monitor (SAM) developed by EGEE for Grid operations
  - Framework to run periodic tests on Grid services to check basic functionalities
  - Computing resources
    - ☞ Can run a job?
    - ☞ Can use the CMS software?
    - ☞ Can read local data?
  - Storage resources
    - ☞ Can copy data in and out?
  - Site is unavailable when fails at least one critical test
  - Availability = fraction of time the site is available



- Site is OK if availability in last 24 hours is
  - > 90% (Tier-1)
  - > 80% (Tier-2)





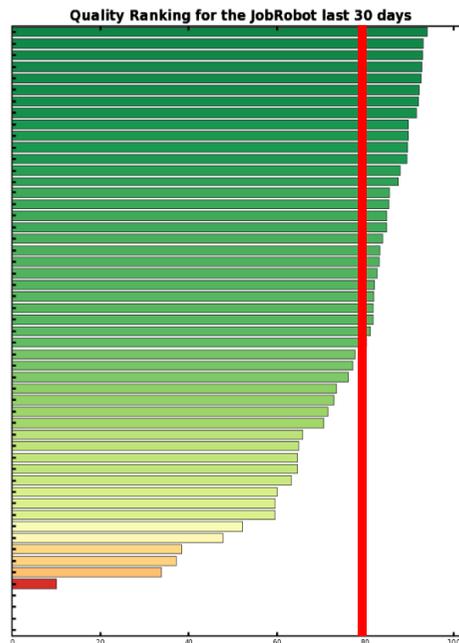
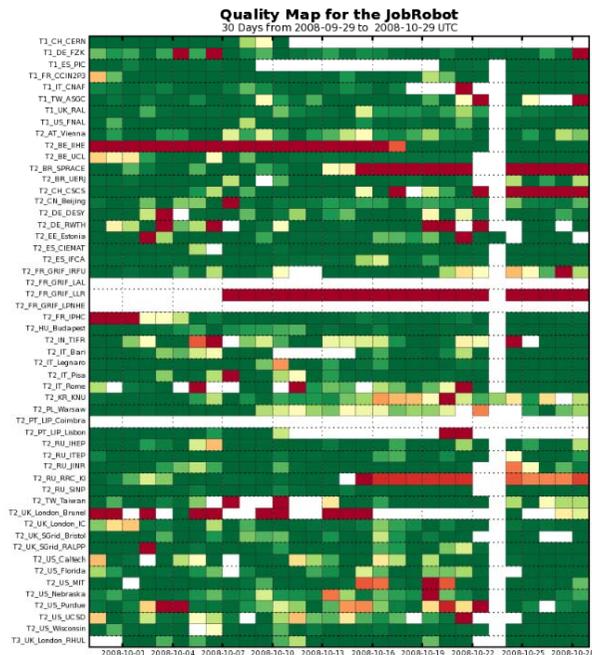
# Job Robot (I)



- A tool to automatically submit fake analysis jobs
  - A set of agents (preparation, submission, collection)
  - Uses CRAB, the CMS analysis job submission tool
  - Reads a dataset at the site of ~500 GB
- Each day the job success rate per site is measured
- Two foreseen operation modes
  - Monitoring mode: constant low rate job submission
  - Commissioning mode: filling site with jobs
    - ☞ To test site under stress
    - ☞ not yet operational

- Summary table
  - Reports no. of failures/successes
  - Access to log files
- Quality plots
  - Vs. Time
  - Site ranking
- Site is OK if success rate
  - > 90% (Tier-1)
  - > 80% (Tier-2)

SITE NAME	SUBMITTED	ABORTED	ERRORS	SUCCESS	EFFICIENCY
T1_ES_PIC	101	0	0	101	100 %
T1_FR_CCIN2P3	202	0	0	202	100 %
T1_IT_CNAF	404	0	0	404	100 %
T1_TW_ASGC	303	0	303	0	0 %
T1_UK_RAL	303	1	9	293	96 %
T1_US_FNAL	202	0	5	111	95 %
T2_AT_Vienna	404	0	0	404	100 %
T2_BE_IHHE	202	0	0	202	100 %
T2_BE_UCL	101	0	0	101	100 %
T2_BR_SPRACE	202	202	0	0	0 %
T2_BR_UERJ	101	0	0	14	100 %
T2_CH_CSCS	101	101	0	0	0 %
T2_CN_Beijing	202	4	0	198	98 %
T2_DE_DESY	101	0	0	101	100 %
T2_ES_CIEMAT	303	0	0	303	100 %
T2_ES_IFCA	303	0	0	303	100 %
T2_FR_GRIF_IRFU	101	0	0	101	100 %
T2_FR_GRIF_LLRF	303	0	303	0	0 %
T2_FR_IPHC	96	0	5	91	94 %
T2_HU_Budapest	404	0	0	404	100 %
T2_IN_TIFR	300	0	0	300	100 %
T2_IT_Bari	101	0	0	101	100 %
T2_IT_Legnaro	303	0	0	303	100 %
T2_IT_Pisa	202	0	0	202	100 %
T2_IT_Rome	101	0	0	101	100 %
T2_PL_Warsaw	202	26	0	176	87 %
T2_RU_IHEP	96	0	0	96	100 %
T2_RU_ITEP	408	4	0	404	99 %
T2_RU_JINR	204	0	0	204	100 %
T2_RU_RRC KI	204	101	79	24	11 %
T2_RU_SINP	408	72	3	333	81 %
T2_TW_Taiwan	202	0	64	138	68 %
T2_UK_London_Brunel	404	0	0	404	100 %
T2_UK_London_IC	202	0	0	202	100 %
T2_UK_SGrid_Bristol	101	0	0	101	100 %
T2_UK_SGrid_RALPP	101	0	0	15	100 %
T2_US_Caltech	404	0	0	404	100 %
T2_US_MIT	101	101	0	0	0 %
T2_US_Nebraska	303	0	41	247	85 %
T2_US_Purdue	101	0	6	95	94 %
T2_US_Wisconsin	202	0	1	201	99 %
T3_UK_London_RHUL	303	0	0	303	100 %
<b>TOTAL (42 sites)</b>	<b>9392</b>	<b>612</b>	<b>819</b>	<b>7687</b>	<b>84 %</b>

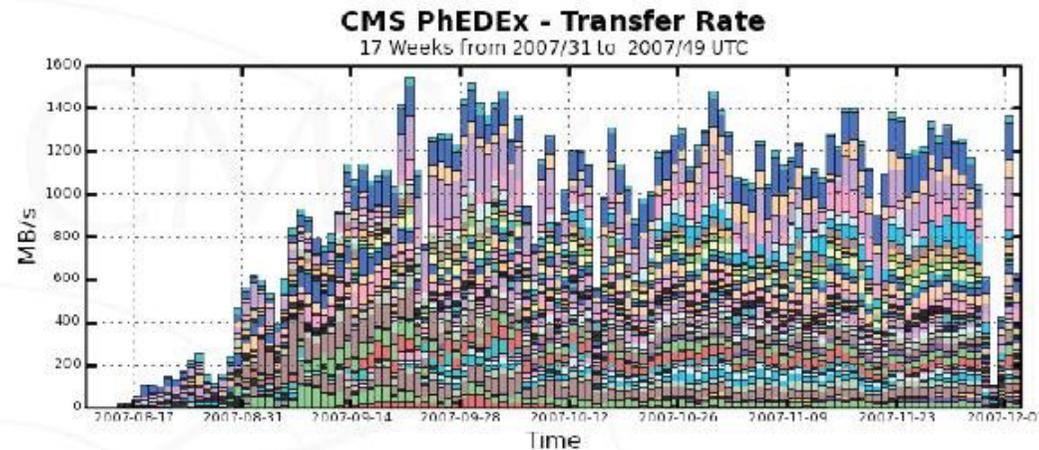




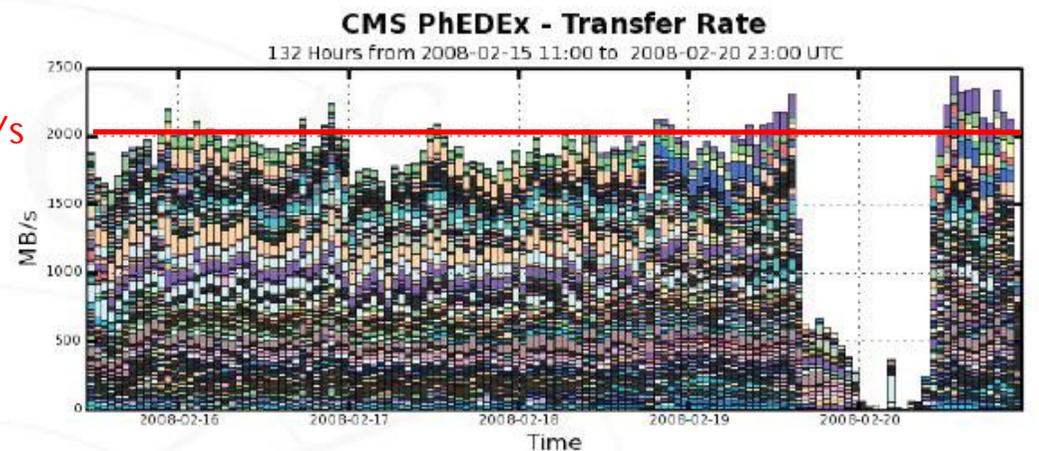
# Link commissioning (I)



- For sites to be usable, these links need to be operational:
  - Tier-0 → Tier-1: to export raw and reconstructed data
  - Tier-1 ↔ Tier-1: for synchronisation and large-scale data reprocessing
  - Tier-1 → Tier-2: for distribution of data to be analyzed at Tier-2 sites
  - Tier-2 → Tier-1: to upload Monte Carlo events generated at Tier-2 sites
- The Debugging Data Transfers (DDT) task force, active since July 2007, defined the metrics, provided a procedure and the tools to test the links and assisted sites in solving problems. Tests were regularly run to verify the status of the links
- The minimum requirements are:
  - 20 MB/s sustained for 24h for Tier-0 → Tier-1 and Tier-1 → Tier-X links
  - 5 MB/s sustained for 24h for Tier-2 → Tier-1 links



**Aug-Dec  
2007**



2 GB/s

**Feb 15-20  
2008 (CCRC'08)**

All 56 Tier-0,1 → Tier-1 links commissioned  
 241 (68%) Tier-1 → Tier-2 links commissioned  
 122 Tier-2 → Tier-1 links commissioned  
 93% of Tier-2 have a working uplink

- The site status is expressed by a single variable
  - **COMMISSIONED**: site is usable and stable
  - **WARNING**: site having a problem right now
  - **UNCOMMISSIONED**: site very unstable in the recent past

### Transition rules

From	U	C	W
To			
U		E for > 2 days over the last 7	E for > 2 days over the last 7
C	O for the last 2 days		O and not E for > 2 days over the last 7
W		E but not E for > 2 days over the last 7	

### Tier-1 sites

Metric	Value
daily SAM availability	≥ 90%
daily JR-MM efficiency	≥ 90%
commissioned downlink with Tier-0	1
commissioned downlinks to Tier-2 sites	≥ 10
commissioned downlinks/uplinks to other Tier-1 sites	≥ 4

### Tier-2 sites

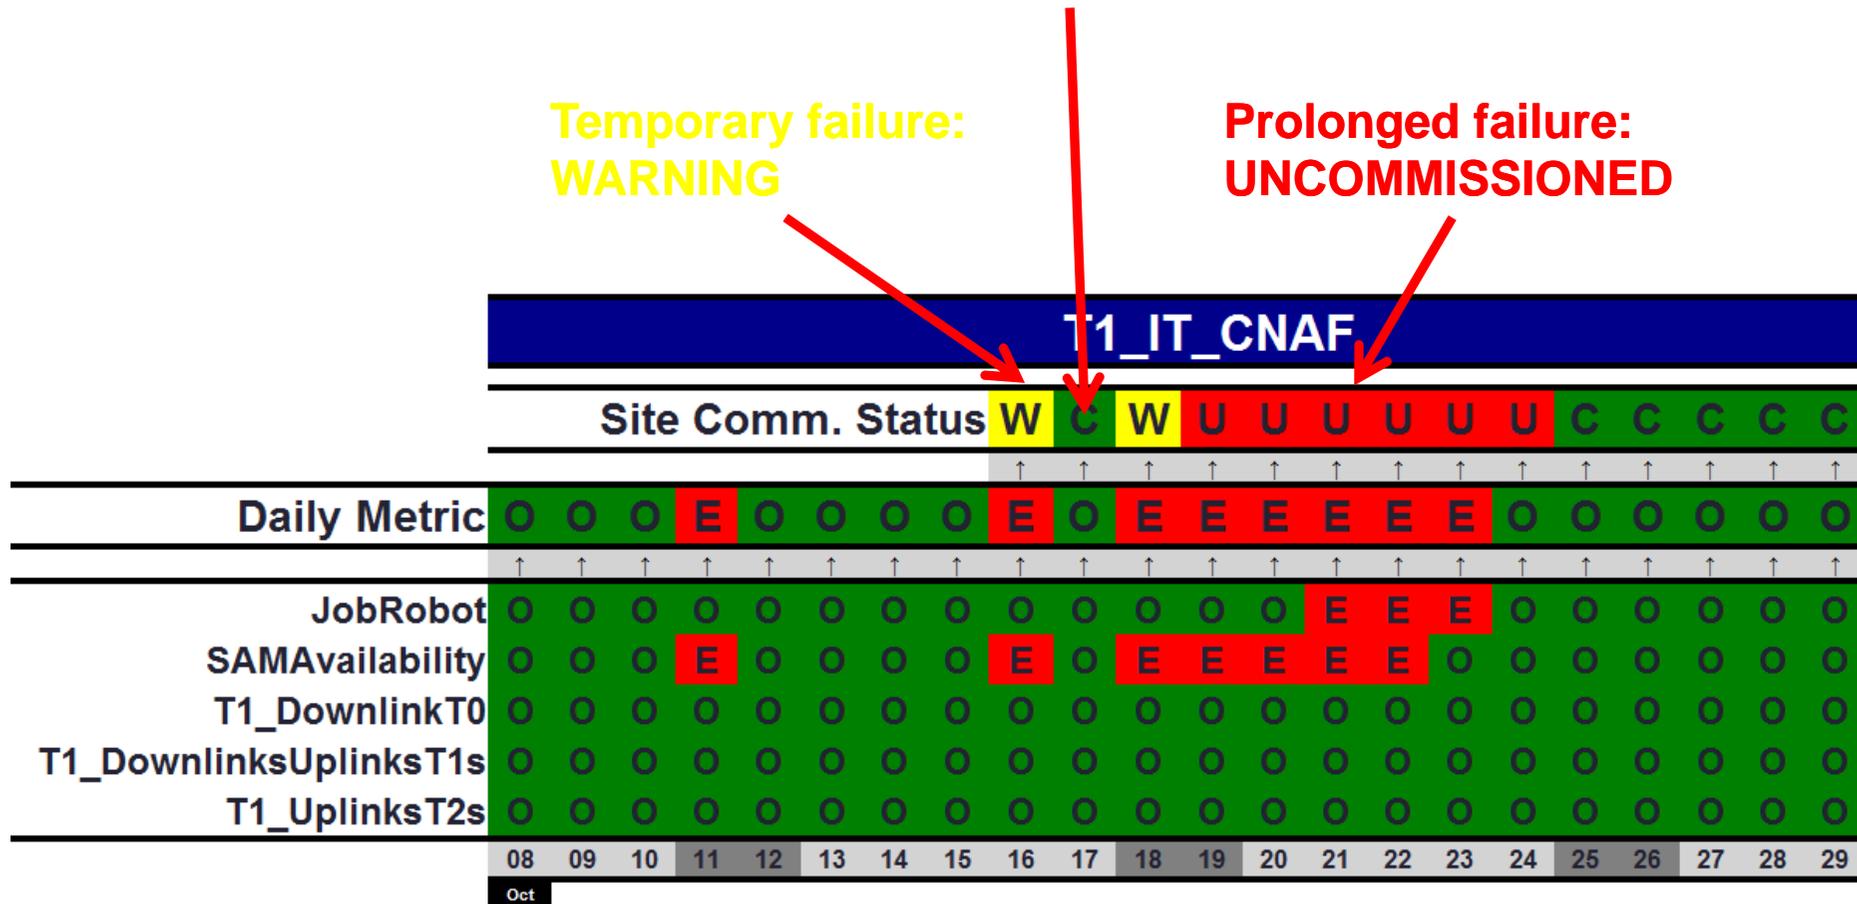
Metric	Value	MC	Analysis
daily SAM availability	≥ 80%	•	•
daily JR-MM efficiency	≥ 80%	•	•
commissioned uplinks with Tier-1 sites	≥ 1	•	
commissioned downlinks with Tier-1 sites	≥ 2		•

- History of last 7 days considered
- Sites should be able to become again commissioned quickly
- Stability > 70% enforced

Few or no recent failures:  
**COMMISSIONED**

Temporary failure:  
**WARNING**

Prolonged failure:  
**UNCOMMISSIONED**





# Site Status Board



Developed by the ARDA Dashboard team and used to visualize all the site commissioning information

Summary table

Site Name	SiteComm JB	Commissioned Links (expand this column)	Site availability	SiteComm Status	Maintenance (expand this column)
T0_CH_CERN	n/a	n/a		n/a	n/a
T1_CH_CERN	n/a	combined		n/a	n/a
T1_DE_FZK	n/a	combined	0	U	n/a
T1_ES_PIC	100%(39)	combined		C	n/a
T1_FR_CCIN2P3	100%(30)	combined		U	n/a
T1_IT_CNAF	100%(50)	combined		C	n/a
T1_TW_ASGC	0%(30)	combined	0	U	n/a
T1_UK_RAL	98%(50)	combined		C	n/a
T1_US_FNAL	97%(30)	combined	1	C	n/a
T2_AT_Vienna	100%(50)	combined	0.20	U	n/a
T2_BE_IHHE	100%(30)	combined	1	C	n/a
T2_BE_UCL	100%(30)	combined		C	n/a
T2_BR_SPRACE	0%(12)	combined	0.00	U	n/a
T2_BR_UERJ	100%(21)	combined	0.30	C	combined
T2_CH_CAF	n/a	combined	n/a	U	n/a
T2_CH_CSCS	0%(47)	combined	0	U	n/a
T2_CN_Beijing	98%(20)	combined	0.00	U	n/a
T2_DE_DESY	100%(20)	combined		C	n/a
T2_DE_RWTH	0%(10)	combined	0.00	U	n/a
T2_EE_Estonia	84%(10)	combined	0.00	C	n/a
T2_ES_CIEMAT	100%(20)	combined	1	C	n/a
T2_ES_IFCA	100%(30)	combined		C	n/a
T2_FI_HIP	n/a	combined	1	n/a	n/a
T2_FR_CCIN2P3	n/a	n/a		n/a	n/a
T2_FR_GRIF_IRFU	27%(30)	combined	0	U	n/a
T2_FR_GRIF_LAL	n/a	combined	n/a	U	n/a
T2_FR_GRIF_IIR	0%(30)	combined	n/a	U	n/a

History plots

Status of SiteComm Status



**Status of Site availability**



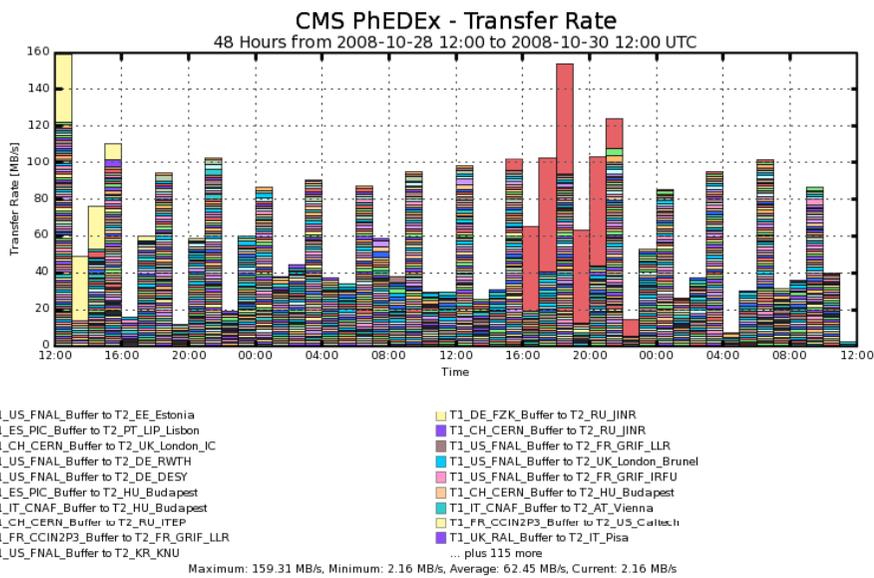
OK: availability > 90% (80%) for Tier-1(2) sites

**Status of SiteComm JR**

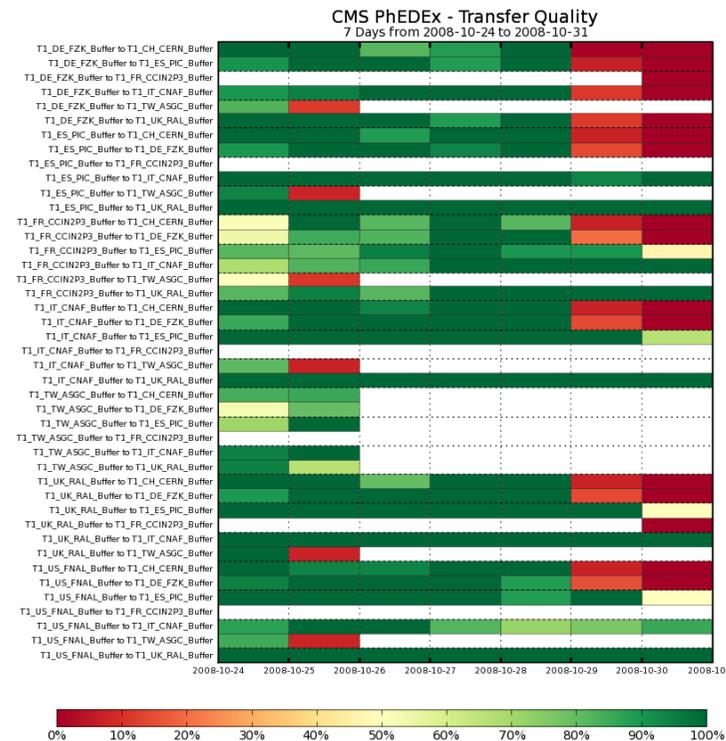


OK: success rate > 90% (80%) for Tier-1(2) sites

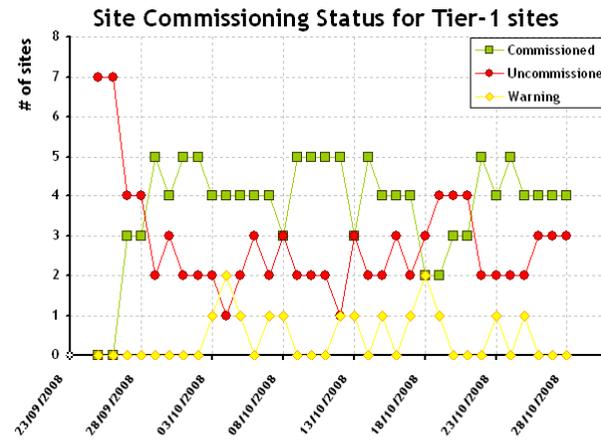
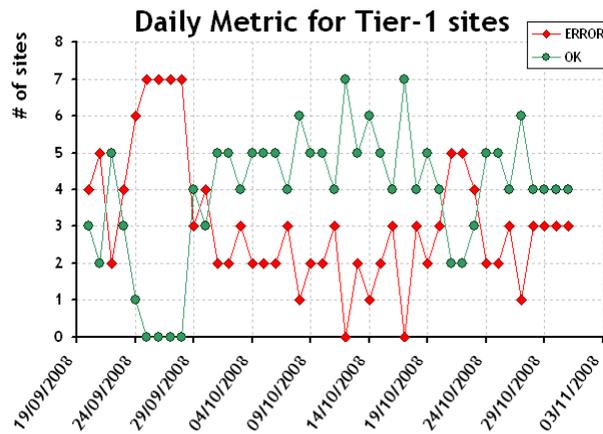
- Count of no. of commissioned links will be replaced with count of links with a good transfer quality
- All links continuously exercised by 0.5 MB/s transfers



Links from Tier-1 to Tier-2



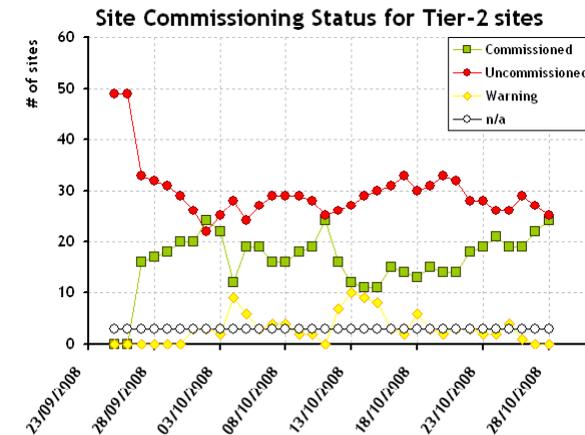
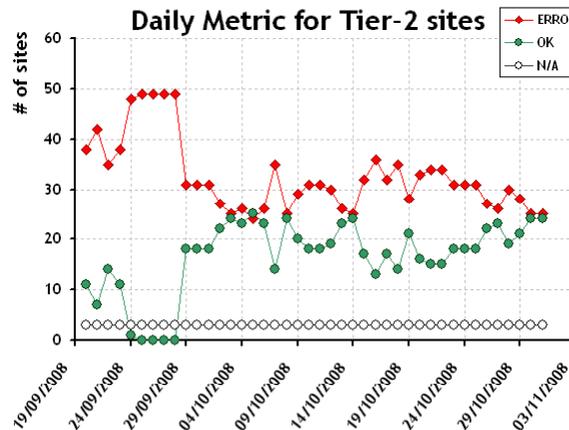
Quality of transfers Tier-1 ↔ Tier-1



Tier-1 sites: usually only 4-5 commissioned at the same time  
 ⇒ to improve

Still a lot of room for improvement

Tier-2 sites: commissioned ~50% of total





# Present and future of the site commissioning



- Machinery almost ready for production use
- Data operations and Physics groups can use the commissioning status to choose the “best” sites
  - Bad Tier-2 sites might be excluded
  - Bad Tier-1 sites are just warned
  - Expected to be a strong incentive at solving problems
- Sites will have an easy way to know if CMS is finding any problems at them
  - Of course they will still be notified of problems via GGUS or Savannah tickets
- The site status will be regularly checked during computing shifts
- Next developments
  - Produce a different status for MC production and analysis
  - Start using the Job Robot in “commissioning mode” to see how sites behave when receiving large numbers of jobs



# Conclusions



- During CCRC'08 all aspects of CMS computing were put under stress, including analysis workflows and Tier-2 sites for the first time
- For regular operations, the site commissioning has provided several tools which can be used
  - By CMS to know which sites are most reliable
  - By sites to know if something is wrong
- The site status is condensed in a single estimator built by continuously probing several functionalities
- The expected result is a significant improvement of the reliability of WLCG sites used by CMS